arXiv:1211.1299v1 [cs.DM] 6 Nov 2012

How applying Myhill-Nerode methods to hypergraphs helps mastering the Art of Trellis Decoding

René van Bevern
Institut für Softwaretechnik und
Theoretische Informatik
Technische Universität Berlin
rene.vanbevern@tu-berlin.de

Michael R. Fellows
School of Engineering and IT
Charles Darwin University
Darwin, Australia
michael.fellows@cdu.edu.au

Serge Gaspers
Computer Science and Engineering
University of New South Wales and
National ICT Australia
sergeg@cse.unsw.edu.au

Frances Rosamond
School of Engineering and IT
Charles Darwin University
Darwin, Australia
frances.rosamond@cdu.edu.au

Abstract

A *trellis* is a graph associated with a linear code that is used for maximum-likelihood decoding. The decoding complexity of a linear code is strongly influenced by the *state complexity* of the trellis, which highly depends on the coordinate permutation of the linear code. The problem of finding the coordinate permutation of a linear code such that the state-complexity of the associated trellis is at most k has been referred to as the *Art of Trellis Decoding* and is NP-hard. We show that it is linear-time solvable if k is constant. To this end, we reduce the problem to HYPERGRAPH CUTWIDTH, which we show to be fixed-parameter linear by, for the first time, applying an analog of the Myhill-Nerode theorem from formal language theory to a hypergraph problem. In contrast, we also exploit Myhill-Nerode methods to show that for instances whose incidence graphs have bounded treewidth, HYPERTREE WIDTH and variants cannot be solved by simple finite tree automata.

1 Introduction

In coding theory, a linear block code is an error-correcting code where each codeword is a vector in a fixed vector space over a fixed finite field such that every linear combination of codewords is also a codeword. After a transmission of data over a noisy or unreliable communication channel, decoding algorithms—such as the Viterbi algorithm [33]—are used to restore the original message. The complexity of many such decoding algorithms heavily depends on the size and structure of a directed graph associated to the code, its minimal *trellis*.

A *trellis* of a linear block code is a directed graph with labeled arcs, one source vertex and one sink vertex. Every directed path from the source vertex to the sink vertex corresponds to the codeword consisting of the sequence of arc-labels encountered on this path, and each codeword of the code corresponds to at least one such path. Moreover, every arc and every vertex of the trellis occur on at least one path from the source to the sink. Since a trellis represents a linear block code, all paths from the source to the sink have the same length. A vertex at distance i from the source is said to have time index i. The state-complexity of a trellis is the maximum number of vertices at any time index. The minimal trellis of a linear block code is the trellis such that any other trellis has at least as many vertices at time index i for every i, $0 \leq i \leq n$. Trellises were introduced by Forney [12] in 1967 in studying the Viterbi algorithm. For an overview of trellis-related decoding algorithms and further applications of trellises, see the survey of Vardy [32].

A generator matrix for a linear block code is a matrix whose rows form a basis of the codeword vector space. If we permute the columns of the generator matrix, we obtain another code, which is equivalent to the original code from an error-correction perspective. However, the permutation may greatly influence the state-complexity of the minimal trellis. The problem of finding a permutation that minimizes the state-complexity of the minimal trellis was termed the "Art of Trellis Decoding" by Massey [27] in 1978. It has attracted a lot of attention in the coding theory literature [32]. The corresponding decision problem, TRELLIS WIDTH, is to decide whether there is a permutation yielding a minimal trellis with state-complexity at most k. It was a long-standing open question whether TRELLIS WIDTH is NP-complete. Forney [15] referred to it as the only significant open problem in the context of trellis decoding of linear block codes. In 2007, TRELLIS WIDTH was finally shown NP-complete by Kashyap [22].

Kashyap [23] also observed that TRELLIS WIDTH is equivalent to MATROID PATHWIDTH. Based on results and conjectures in Matroid Minor theory [18], Kashyap expressed a strong belief that TRELLIS WIDTH is polynomial-time solvable if k is a constant. Indeed, the class of matroids with pathwidth at most k is closed under taking matroid minors [23], has bounded branchwidth, and membership in a minor-closed matroid family can be tested in cubic time for matroids with bounded branchwidth [21]. This asserts that, for every constant k, there exists a cubic-time algorithm for Trellis Width.[1] However, the list of minimal forbidden minors and, hence, the algorithm, is only known explicitly for $k = 1$ [23]. For $k > 1$, the algorithm, based on well-quasiordering, is nonuniformly fixed-parameter tractable (there is a different algorithm for each k) [7], and nonconstructive: for each constant k, we only know that an $O(n^3)$-time algorithm exists, but it is unknown what it is, and what kind of dependence on k is hidden in the big-O.

Our Contributions. We give an explicit algorithm for Trellis Width that, for every fixed constant k, runs in linear time. Stated in terms of Parameterized Complexity [7], our algorithm is *fixed-parameter linear.*

As a first step, we first use structural properties of minimal trellises to reduce Trellis Width to Hypergraph Cutwidth. To show that Hypergraph Cutwidth is fixed-parameter linear, we extend the graph-theoretic analog [8] of the Myhill-Nerode characterization of regular languages first to colored graphs, and then to hypergraphs. Our hypergraph analog of the Myhill-Nerode theorem can then be used to derive fixed-parameter linear algorithms for hypergraph problems on instances with bounded *incidence treewidth*, i.e., instances whose incidence graphs have bounded treewidth. Incidence treewidth is a particularly nice parameter since it is bounded from above by other well-known width measures of hypergraphs: for example, by the treewidth of the primal graph plus one [24], and by the treewidth of the dual graph plus one [30].

Our fixed-parameter linear algorithm for Hypergraph Cutwidth improves the $O(n^{k^2+3k+3})$ time algorithm of Miller and Sudborough [28]. Besides Hypergraph Cutwidth, Trellis Width, and Matroid Pathwidth, our extensions of the Myhill-Nerode theorem to colored graphs and hypergraphs are likely to apply to other problems. Colored or annotated graphs allow for more realism in problem modeling and often arise as subproblems when solving pure graph problems. It is also straightforward to use our methods for annotated hypergraphs.

The original Myhill-Nerode theorem can be used both positively and negatively: to show that a language is regular, and to show that a language is not regular. As an example of a negative use of our hypergraph Myhill-Nerode analog, we prove that for instances with bounded incidence treewidth, the Hypertree Width, Generalized Hypertree Width, and Fractional Hypertree Width problems cannot be solved by bottom-up finite-state tree automata (shortly, finite tree automata) whose inputs are parse trees corresponding to bounded-width tree decompositions of the incidence graph. Moreover, the proof leads us to the conjecture that these problems are not *fixed-parameter tractable* with respect to the parameter incidence treewidth t, that is, we conjecture that there are no algorithms for these problems running in time $f(t) \cdot n^c$, where n is the input size, c is a constant and f only depends on t.

2 Preliminaries

For standard graph-theoretic notions, we follow the definitions and notations of Diestel [6].

Graph Decompositions. A *tree decomposition* of a graph $G = (V, E)$ is a pair $(\{X_i : i \in I\}, T)$ where $X_i \subseteq V$, $i \in I$, are called *bags* and T is a tree with elements of I as nodes such that:

1. for each edge $\{u, v\} \in E$, there is an $i \in I$ such that $\{u, v\} \subseteq X_i$, and
2. for each vertex $v \in V$, $T[\{i \in I : v \in X_i\}]$ is a (connected) tree with at least one node.

The *width* of a tree decomposition is $\max_{i \in I} |X_i| - 1$. The *treewidth* of G is the minimum width taken over all tree decompositions of G. The notions of *path decomposition* and *pathwidth* of G are defined the same way, except that T is restricted to be a path.

Hypergraphs. A *hypergraph* H is a couple (V, E), where V is a set of *vertices* and E a multiset of *hyperedges* such that $e \subseteq V$ for each $e \in E$. Let $H = (V, E)$ be a hypergraph. The *primal graph* of H, denoted $\mathcal{G}(H)$, is the graph with vertex set V that has an edge $\{u, v\}$ if there exists a hyperedge in H incident to both u and v. It is sometimes called the Gaifman graph of H. The *incidence graph* of H, denoted $\mathcal{I}(H)$, is the bipartite graph (V', E') with vertex set $V' = V \cup E$ and for $v \in V$ and $e \in E$, there is an edge $\{v, e\} \in E'$ if $v \in e$.

Hypergraph decompositions. Let H be a hypergraph. The *treewidth* of H is the treewidth of $\mathcal{G}(H)$. The *incidence treewidth* of H is the treewidth of $\mathcal{I}(H)$.

Generalized hypertree width is defined with respect to tree decompositions of $\mathcal{G}(H)$, however, the width of the tree decompositions is measured differently. Suppose H has no isolated vertices (otherwise, remove them). A *cover* of a bag is a set of hyperedges such that each vertex in the bag is contained in at least one of these hyperedges. The *cover width* of a bag is the minimum number of hyperedges covering it. The *cover width* of a tree decomposition is the maximum cover width of any bag in the decomposition. The *generalized hypertree width* of H is the minimum cover width over all tree decompositions of $\mathcal{G}(H)$.

The *hypertree width* of H is defined in a similar way, except that, additionally, the tree of the decomposition is rooted and a hyperedge e can only be used in the cover of a bag X_i if X_i contains all vertices of e that occur in bags of the subtree rooted at the node i.

The *fractional hypertree width* of H is defined in a similar way as the generalized hypertree width, except that it uses fractional covers: in a *fractional cover* of a bag, each hyperedge is assigned a non-negative weight, and for each vertex in the bag, the sum of the weights of the hyperedges incident to it is at least 1. The *fractional cover width* of the bag is the minimum total sum of all hyperedges of a fractional cover.

3 Trellis Width as a Hypergraph Problem

In this section we formally define the Trellis Width and the Hypergraph Cutwidth problems, and we give a parameterized reduction from the former to the latter. Then we briefly consider the related TRELLIS BRANCH COMPLEXITY problem and observe that it is fixed-parameter linear.[2]

---

[1] This has also been observed by Navin Kashyap [private communication].

Trellis Width. Let q be a prime power. We consider linear (block) codes over the finite field $\mathbb{F}_q$. A *linear code* of length n and rank d is a linear subspace C with dimension d of the vector space $\mathbb{F}_q^n$. The vectors in C are called *codewords*. Linear codes are usually represented by a *generator matrix* whose d rows are a basis of C.

A *trellis* $D = (V, A)$ for a linear code C of length n is an arc-labeled directed acyclic graph with one *source* node s and one *sink* node t with arc labels from $\mathbb{F}_q$ such that

1. every vertex and every arc of the trellis is on at least one s-t-path, and
2. the set of n-tuples obtained by reading off the arc labels obtained in traversing all s-t-paths is C.

The number of arcs that need to be traversed to reach a vertex v from s is the *time index* associated with v. We denote by $V_D(i)$ the set of vertices at time index i in D. Thus, $V_D(0) = \{s\}$ and $V_D(n) = \{t\}$.

It has been shown [13, 16, 25, 29] that every linear code C has a unique *minimal* trellis, characterized by the fact that any other trellis for C has at least as many vertices at time index i, for every i, $0 \leq i \leq n$, where n is the length of the code. In his survey, Vardy [32] presents four different constructions of the minimal trellis that appeared in the literature [2, 13, 25, 27]. From these constructions, it follows that the number of vertices at each time index is a power of q. The *state-complexity* or *trellis width* of a code is $\max_{0 \leq i \leq n}\{s_i\}$, where $q^{s_i} = |V_D(i)|$ and D is its minimal trellis.

The decoding complexity—for example, by the Viterbi algorithm [33]—is closely tied to the width of the associated trellis. However, a permutation of the columns of a code may greatly influence the width of the resulting minimal trellis, whereas from an error-correction viewpoint the performance of the code remains the same. This leads us to the following computational problem.

> TRELLIS WIDTH
> *Input:* A prime power q, a generator matrix A of a linear block code over $\mathbb{F}_q$, and a natural number k.
> *Question:* Is there a column permutation of A that yields a code with state-complexity at most k?

The problem was only recently shown to be NP-hard by Kashyap [22, 23], settling a long-standing conjecture. It is also equivalent to the MATROID PATHWIDTH problem [23, 32].

Hypergraph Cutwidth. Let $H = (V, E)$ be a hypergraph. A *linear layout* of H is an injective map $l\colon V \to \mathbb{R}$ of vertices onto the real line. The *cut at position* $i \in \mathbb{R}$ with respect to l, denoted $\theta_{l,H}(i)$, is the set of hyperedges that contain at least two vertices v, w such that $l(v) < i < l(w)$. The *cutwidth* of the layout l is $\max_{i\in\mathbb{R}} |\theta_{l,H}(i)|$. The *cutwidth* of the hypergraph H is the minimum cutwidth over all the linear layouts of H. The HYPERGRAPH CUTWIDTH problem is then defined as follows.

> HYPERGRAPH CUTWIDTH
> *Input:* A hypergraph $H = (V, E)$ and a natural number k.
> *Question:* Does H have cutwidth at most k?

The problem naturally generalizes the NP-complete [17] GRAPH CUTWIDTH problem for which several fixed-parameter algorithms are known [1, 4, 9, 10, 31]. Cahoon and Sahni [5] designed algorithms for HYPERGRAPH CUTWIDTH with $k \leq 2$, with running time $O(n)$ for $k = 1$ and running time $O(n^3)$ for $k = 2$, where n is the number of vertices. For arbitrary k, Miller and Sudborough [28] designed an algorithm with running time $O(n^{k^2+3k+3})$.

Trellis Width as Hypergraph Cutwidth. Exploiting the structure of minimal trellises, one can reduce TRELLIS WIDTH to HYPERGRAPH CUTWIDTH. Let D be the minimal trellis for a linear code C over $\mathbb{F}_q$ with generator matrix A. Let n and d denote the number of columns and rows of A, respectively. Since the rows of A are a basis of C, no row consists of zeros only. For a row $c = (c_1, \ldots, c_n)$ of A, the *span* of c, denoted $s(c)$, is the interval $[a, b]$ where a is the smallest integer i for which c_i is nonzero and b is the largest integer i for which c_i is nonzero. For any $i, 1 \leq i \leq n$, denote by p_i the number of rows $c = (c_1, \ldots, c_n)$ of A with $s(c) \subseteq [1, i]$, and denote by f_i the number of rows $c = (c_1, \ldots, c_n)$ of A with $s(c) \subseteq [i + 1, n]$. By the State Space Theorem of Forney and Trott [16], we have

State Space Theorem ([16]). *$|V_D(i)| = q^{d-p_i-f_i}$ for $i = 0, 1, \ldots, n$.*

We are now ready to present a parameterized reduction [7].

Lemma 1. *There is a linear-time parameterized reduction from TRELLIS WIDTH to HYPERGRAPH CUTWIDTH.*

Proof. Let (q, A, k) be an instance for TRELLIS WIDTH. We construct an equivalent instance $(H = (V, E), k)$ for HYPERGRAPH CUTWIDTH as follows: We set $V := \{v_1, \ldots, v_n\}$, where n is the number of columns of A. For each row $(c_1, \ldots, c_n)$ of A, let there be a hyperedge $e \in E$ such that $v_i \in e \iff c_i \neq 0$. This describes all the hyperedges of H. Note that the parameter “trellis width” of the input instance equals the parameter “cutwidth” of the output instance and that this construction works in linear time.

Suppose there is a column permutation π of A that yields a code C with state-complexity at most k. Consider the linear layout l that maps v_i to $\pi(i)$, $1 \leq i \leq n$. The cutwidth of l is $\max_{i\in\mathbb{R}} |\theta_{l,H}(i)| = \max_{i\in\{1,\ldots,n-1\}} |\theta_{l,H}(i + 0.5)|$. Let $s_{i,\pi}$ denote the number of codewords that have a non-zero entry in the first i columns and a non-zero entry in the last $n - i$ columns of the matrix obtained from A by the permutation π of the columns. Consider an arbitrary $i \in \{1, \ldots, n-1\}$. Then, $|\theta_{l,H}(i+0.5)| = s_{i,\pi}$. By the State Space Theorem, the minimal trellis of C has $q^{s_{i,\pi}}$ vertices at time index i. Therefore, $|\theta_{l,H}(i + 0.5)| = s_{i,\pi} \leq k$. Since i was arbitrary, it follows that the cutwidth of H is at most k.

Suppose there is a linear layout l of H with cutwidth at most k. Without loss of generality, assume l is a bijection from V to $\{1, \ldots, n\}$. Consider the column permutation π

---

[2] When we do not specify a parameter, we implicitly assume the *standard parameterization*, i.e., the parameter is what one would optimize in the optimization version of the problem.

such that $l(v_i) = \pi(i)$. Let C denote the linear code whose generator matrix is the column permutation π of A. Again, it follows from the State Space Theorem, that the minimal trellis for C has state-complexity at most k. ☐

By this reduction, any fixed-parameter linear algorithm for HYPERGRAPH CUTWIDTH translates into a fixed-parameter linear algorithm for TRELLIS WIDTH and, therefore, also for MATROID PATHWIDTH.

Trellis Branch Complexity. Besides minimizing the maximum vertex-count in the minimal trellis at any time index, the problem to minimize the arc-count between any two consecutive time indices has also been proposed in the literature [14]. Denote by $E_D(i)$ the set of arcs that end at a vertex from $V_D(i)$. The *branch-complexity* of a code is $\max_{1 \leq i \leq n}\{b_i\}$ where $q^{b_i} = |E_D(i)|$ and D is its minimal trellis.

TRELLIS BRANCH COMPLEXITY
Input: A prime power q, a generator matrix A of a linear block code over $\mathbb{F}_q$, and a natural number k.
Question: Is there a column permutation of A that yields a code with branch-complexity at most k?

Similarly to the State Space Theorem, the number of arcs between two consecutive time indices is easily read off the generator matrix.

Branch Space Theorem ([14]).
$|E_D(i)| = q^{d-p_{i-1}-f_i}$ *for* $i = 1, 2, \ldots, n$.

Replacing each non-zero entry in the generator matrix by a '1' easily reduces TRELLIS BRANCH COMPLEXITY to GATE MATRIX LAYOUT.

GATE MATRIX LAYOUT
Input: A Boolean matrix M and an integer k.
Question: Is there a permutation of the columns of M so that, if in each row we enforce the consecutive-ones property, then no column contains more than k ones?

Here, *enforcing the consecutive-ones property* in a row means replacing by a '1' each '0' that occurs between two '1's. The GATE MATRIX LAYOUT problem with parameter k has a parameterized reduction to PATHWIDTH [10]. Since PATHWIDTH is fixed-parameter linear [3] and all reductions are linear-time, we have the following theorem.

Theorem 1. TRELLIS BRANCH COMPLEXITY *is fixed-parameter linear.*

4 Myhill-Nerode Analog for Colored Graphs and Hypergraphs

The aim of this section is to generalize the Myhill-Nerode analog for graphs [8] to colored graphs and hypergraphs. This will give us the tools to show that HYPERGRAPH CUTWIDTH is fixed-parameter linear.

The methodology is first extended to colored graphs. From this, we obtain a Myhill-Nerode analog for hypergraphs, since every hypergraph can be represented as its incidence graph with two vertex types: vertices that represent hyperedges and vertices that represent hypergraph-vertices.

In the last part of the section, we finally describe how our Myhill-Nerode analog yields linear-time algorithms for hypergraph problems.

We follow and adapt the notation used by Downey and Fellows [7, Section 6.4].

4.1 Colored Graphs

We now develop an analog of the Myhill-Nerode theorem for colored graphs. The original Myhill-Nerode theorem is stated for languages in terms of concatenations of words. Hence, we clarify what concatenating colored graphs means.

Definition 1. A *t-boundaried graph G* is a graph with t distinguished vertices that are labeled from 1 to t. These labeled vertices are called *boundary vertices* and the *boundary*, $\partial(G)$, denotes the set of boundary vertices of G.

Let G_1 and G_2 be t-boundaried graphs whose vertices are colored with colors from $\{1, \ldots, c_{\max}\}$. We say that G_1 and G_2 are *color-compatible* if the vertices with the same labels in $\partial(G_1)$ and $\partial(G_2)$ have the same color.

For two color-compatible t-boundaried graphs, we denote by $G_1 \oplus_c G_2$ the colored graph obtained by taking the disjoint union of G_1 and G_2 and identifying each vertex of $\partial(G_1)$ with the vertex of $\partial(G_2)$ with the same label, wherein vertex colors are inherited from G_1 and G_2.

Let $\mathcal{U}^{\text{large}}_{t,c_{\max}}$ be the universe of $\{1, \ldots, c_{\max}\}$-colored t-boundaried graphs and $F \subseteq \mathcal{U}^{\text{large}}_{t,c_{\max}}$. We define the *canonical right congruence* $\sim_F$ for F as follows: for $G_1, G_2 \in \mathcal{U}^{\text{large}}_{t,c_{\max}}$, $G_1 \sim_F G_2$ if and only if G_1 and G_2 are color-compatible and for all color-compatible $H \in \mathcal{U}^{\text{large}}_{t,c_{\max}}$, it holds that $G_1 \oplus_c H \in F \iff G_2 \oplus_c H \in F$.

The *index* of $\sim_F$ is its number of equivalence classes.

Definition 1 is illustrated in Figure 1. Two colored t-boundaried graphs G_1 and G_2 are *isomorphic* and we write $G_1 \cong G_2$ if there is a graph isomorphism for the underlying (ordinary) graphs mapping each vertex to a vertex with the same color.

Before we can state our analog of the Myhill-Nerode theorem for colored graphs, we show that every colored graph of treewidth at most t can be generated using a constant number of graph operations. To this end, we use the following set of operators. For generating graphs of only one color, the given operators coincide with those given by Downey and Fellows [7, Section 6.4].

Definition 2. The *size-$(t+1)$ parsing operators* for $\{1, \ldots, c_{\max}\}$-colored graphs are:

i) $\{\emptyset_{n_1,\ldots,n_{c_{\max}}} : \sum_{i=1}^{c_{\max}} n_i = t+1\}$ is a family of nullary operators that creates boundary vertices $1, \ldots, t+1$, of which the first n_1 vertices get color 1, the next n_2 vertices get color 2, and so on.
ii) γ is a unary operator that cyclically shifts the boundary. That is, γ moves label j to the vertex with label $j+1 \pmod{t+1}$.
iii) i is a unary operator that assigns the label 1 to the vertex currently labeled 2 and label 2 to the vertex with label 1.
iv) e is a unary operator that adds an edge between the vertices labeled 1 and 2.

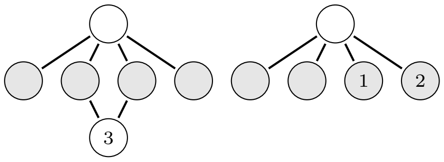

(a) A 2-colored 3-boundaried graph G.

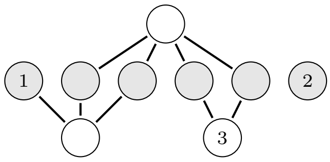

(b) A 2-colored 3-boundaried graph H.

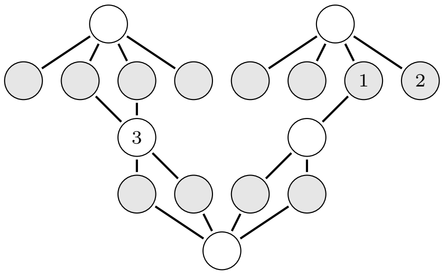

(c) The glued graph $G \oplus_{\mathrm{c}} H$.

Figure 1: Two color-compatible 3-boundaried graphs G and H and their glued graph, where the boundary vertices are marked by their label.

v) $\{u_\ell : 1 \leq \ell \leq c_{\max}\}$ is a family of unary operators that add a new vertex of color ℓ and label it 1, unlabeling the vertex previously labeled 1.

vi) $\oplus_{\mathrm{c}}$ is our gluing operator from Definition 1.

Obviously, for a constant number of colors $c_{\max}$, the set of size-$(t+1)$ parsing operators is finite. Following the proof of Downey and Fellows [7, Theorem 6.72], it is easy to verify that the graphs generated by the operators in Definition 2 are isomorphic precisely to the graphs of treewidth at most t and at least $t+1$ vertices. The same proof shows how, from a width-t tree decomposition of a colored graph, a linear-size parse tree for an isomorphic graph can be obtained in linear time.

Definition 3. The set $\mathcal{U}^{\mathrm{small}}_{t,c_{\max}}$ denotes the set of $\{1, \ldots, c_{\max}\}$-colored t-boundaried graphs that can be generated by the operators in Definition 2.

Theorem 2. *Let $F \subseteq \mathcal{U}^{\mathrm{small}}_{t,c_{\max}}$ be a family of graphs. The following statements are equivalent:*

i) The parse trees corresponding to graphs in F are recognizable by a finite tree automaton.

ii) The canonical right congruence $\sim_F$ has finite index over $\mathcal{U}^{\mathrm{small}}_{t,c_{\max}}$.

Our proof of Theorem 2 adapts the proof of Downey and Fellows [7, Theorem 6.77]. To this end, we only need to show that our operators satisfy the Parsing Replacement Property.

Definition 4. *(Parsing Replacement Property).* An n-ary operator $\otimes$ has the *Parsing Replacement Property* if, for the arguments $H_1, \ldots, H_n \in \mathcal{U}^{\mathrm{small}}_{t,c_{\max}}$ of $\otimes$ and each i, there is a graph $G \in \mathcal{U}^{\mathrm{small}}_{t,c_{\max}}$ such that $H_1 \otimes \cdots \otimes H_n \cong H_i \oplus_{\mathrm{c}} G$.

Lemma 2. The size-$(t+1)$ parsing operators in Definition 2 have the Parsing Replacement Property.

Proof. Let $H \in \mathcal{U}^{\mathrm{small}}_{t,c_{\max}}$ and let $\emptyset^*_H$ denote the graph that contains only $t+1$ boundary vertices, no edges, and that is color-compatible with H. Finally, let $\emptyset^\ell_H$ denote the graph that contains $t+1$ boundary vertices and one additional vertex with color ℓ and that is color-compatible with H. Note that $\emptyset^\ell_H, \emptyset^*_H \in \mathcal{U}^{\mathrm{small}}_{t,c_{\max}}$: we can generate them using the operators in Definition 2 because $\emptyset^*_H$ and $\emptyset^\ell_H$ have at least $t+1$ vertices and treewidth at most t. Now, the lemma immediately follows, noting that for each graph $H \in \mathcal{U}^{\mathrm{small}}_{t,c_{\max}}$,

$$\begin{aligned} \gamma H \cong iH &\cong H \oplus_{\mathrm{c}} \emptyset^*_H, \text{ since } \cong \text{ ignores boundary labels,} \\ eH &\cong H \oplus_{\mathrm{c}} e(\emptyset^*_H), \\ u_\ell H &\cong H \oplus_{\mathrm{c}} \emptyset^\ell_H, \end{aligned}$$

and that $\oplus_{\mathrm{c}}$ and $\emptyset_{n_1,\ldots,n_{c_{\max}}}$ trivially have the Parsing Replacement Property. ☐

Having proven Lemma 2, the proof of Theorem 6.77 in the book of Downey and Fellows [7] also proves Theorem 2 and we have a handy tool to generalize the Myhill-Nerode theorem to hypergraphs.

4.2 Lifting to Hypergraphs

To make the method accessible to hypergraph problems we lift the Myhill-Nerode theorem for colored graphs of the previous subsection to hypergraphs.

Definition 5. A t-boundaried hypergraph G is a hypergraph with t distinguished vertices and hyperedges labeled from 1 to t. Two t-boundaried hypergraphs are *gluable* if no vertex of one hypergraph has the label of a hyperedge of the other hypergraph.

Let G_1 and G_2 be gluable t-boundaried hypergraphs. We denote by $G_1 \oplus_{\mathrm{h}} G_2$ the t-boundaried hypergraph obtained by taking the disjoint union of G_1 and G_2, identifying each labeled vertex of G_1 with the vertex of G_2 with the same label, and replacing the hyperedges with label ℓ by the union of these hyperedges.

Let $\mathcal{H}^{\mathrm{large}}_t$ be the universe of t-boundaried hypergraphs and $F \subseteq \mathcal{H}^{\mathrm{large}}_t$. We define the *canonical right congruence* $\sim_F$ for F as follows: for $G_1, G_2 \in \mathcal{H}^{\mathrm{large}}_t$, $G_1 \sim_F G_2$ if and only if for all $H \in \mathcal{H}^{\mathrm{large}}_t$, $G_1 \oplus_{\mathrm{h}} H \in F \iff G_2 \oplus_{\mathrm{h}} H \in F$.

To prove a Myhill-Nerode theorem for hypergraphs, we still need a way to create hypergraphs from parsing operators, as we did using the operators from Definition 2 for colored graphs. To this end, we indeed simply use the operators from Definition 2, observing that every bipartite graph can be interpreted as the incidence graph of a hypergraph. Moreover, if the two disjoint independent sets of a two-colored bipartite graph have distinct colors, then we can interpret this bipartite graph as a hypergraph in a unique way.

Definition 6. A *well-colored t-boundaried graph* is a $\{1, 2\}$-colored t-boundaried graph $G = (U \uplus W, E)$, where the vertices in U have color 1, the vertices in W have color 2, and where U and W are independent sets.

For a well-colored t-boundaried graph $G = (U \uplus W, E)$, we denote by $\mathcal{H}(G)$ the t-boundaried hypergraph with the vertex set U and the edge set $\{N(w) : w \in W\}$. Moreover, vertices of $\mathcal{H}(G)$ inherit their label from G and edges $e = N(w), w \in W$ of $\mathcal{H}(G)$ inherit the label of w. For a set $U \subseteq \mathcal{U}_{t,2}^{\text{large}}$, we denote $\mathcal{H}(U) := \{H \in \mathcal{H}_t^{\text{large}} : H = \mathcal{H}(G), G \in U\}$.

The *incidence graph* of $\mathcal{H}(G)$ is G and the *incidence treewidth* of $\mathcal{H}(G)$ is the treewidth of G.

Obviously, every $H \in \mathcal{H}_t^{\text{large}}$ is $\mathcal{H}(G)$ for some $G \in \mathcal{U}_{t,2}^{\text{large}}$. Also, every set $\mathcal{F} \subseteq \mathcal{H}_t^{\text{large}}$ is $\mathcal{F} = \mathcal{H}(F)$ for some $F \subseteq \mathcal{U}_{t,2}^{\text{large}}$. Definitions 5 and 6 are illustrated in Figure 2.

Theorem 3. *Let $\mathcal{F} \subseteq \mathcal{U}_{t,2}^{small}$ be a set of well-colored t-boundaried graphs. The following statements are equivalent:*

i) The parse trees corresponding to graphs in F are recognizable by a finite tree automaton.

ii) The canonical right congruence $\sim_{\mathcal{H}(F)}$ has finite index over $\mathcal{H}(\mathcal{U}_{t,2}^{small})$.

Proof. Let $\bar{F} := \{G \in \mathcal{U}_{t,2}^{\text{small}} : \mathcal{H}(G) \in \mathcal{H}(F)\}$, that is, for each $G \in F$, the set $\bar{F}$ contains additionally all t-boundaried graphs $G' \in \mathcal{U}_{t,2}^{\text{small}}$ for which $H(G) = H(G')$. Clearly, every such G' is isomorphic to G, that is, $\bar{F}$ contains F and t-boundaried graphs isomorphic to graphs in F. Hence, the parse trees corresponding to t-boundaried graphs in F are the same as those corresponding to $\bar{F}$. Moreover, $\mathcal{H}(F) = \mathcal{H}(\bar{F})$. Hence, it remains to show that $\sim_{\mathcal{H}(\bar{F})}$ has finite index over $\mathcal{H}(\mathcal{U}_{t,2}^{\text{small}})$ if and only if $\sim_{\bar{F}}$ has finite index over $\mathcal{U}_{t,2}^{\text{small}}$.

Assume that $\sim_{\bar{F}}$ has infinite index over $\mathcal{U}_{t,2}^{\text{small}}$. Then there is a family $G_1, G_2, G_3, \ldots$ of graphs in $\mathcal{U}_{t,2}^{\text{small}}$ that are pairwise inequivalent under $\sim_{\bar{F}}$. Since there are only a finite number of possibilities to assign two colors to t boundary vertices, there are an infinite number of color-compatible graphs among $G_1, G_2, \ldots$. Moreover, notice that all G_i that are not well-colored t-boundaried graphs are equivalent under $\sim_{\bar{F}}$ (they cannot be completed to graphs in $\bar{F}$ by gluing any graph onto them). Therefore, without loss of generality, we assume that $G_1, G_2, \ldots$ are pairwise color-compatible well-colored t-boundaried graphs. Now, for each pair G_i, G_j, there is a graph $H_{i,j} \in \mathcal{U}_{t,2}^{\text{small}}$ such that $G_i \oplus_{\text{c}} H_{i,j} \in \bar{F}$ but $G_j \oplus_{\text{c}} H_{i,j} \notin \bar{F}$. From $G_i \oplus_{\text{c}} H_{i,j} \in \bar{F}$, it follows that $H_{i,j}$ is a well-colored t-boundaried graph that is color-compatible with G_i. Now, we have $\mathcal{H}(G_i) \oplus_{\text{h}} \mathcal{H}(H_{i,j}) = \mathcal{H}(G_i \oplus_{\text{c}} H_{i,j}) \in \mathcal{H}(\bar{F})$. Moreover, $\mathcal{H}(G_j) \oplus_{\text{h}} \mathcal{H}(H_{i,j}) = \mathcal{H}(G_j \oplus_{\text{c}} H_{i,j}) \notin \mathcal{H}(\bar{F})$. That is, $\mathcal{H}(G_i) \nsim_{\mathcal{H}(\bar{F})} \mathcal{H}(G_j)$ and therefore $\sim_{\mathcal{H}(\bar{F})}$ has infinite index.

Assume that $\sim_{\mathcal{H}(\bar{F})}$ has infinite index over $\mathcal{H}(\mathcal{U}_{t,2}^{\text{small}})$. Then, there is a family $\mathcal{H}(G_1), \mathcal{H}(G_2), \mathcal{H}(G_3), \ldots$ of hypergraphs in $\mathcal{H}(\mathcal{U}_{t,2}^{\text{small}})$ that are pairwise inequivalent under $\sim_{\bar{F}}$. Since there are only a finite number of partitions of t labels into hyperedge-labels and vertex-labels, there is an infinite number of pairwise gluable hypergraphs among $\mathcal{H}(G_1), \mathcal{H}(G_2), \ldots$. Therefore, without loss of generality, assume that all these hypergraphs are pairwise gluable. Now, for each pair $\mathcal{H}(G_i)$, $\mathcal{H}(G_j)$, there is a hypergraph $\mathcal{H}(H_{i,j}) \in \mathcal{H}(\mathcal{U}_{t,2}^{\text{small}})$ such that $\mathcal{H}(G_i) \oplus_{\text{h}} \mathcal{H}(H_{i,j}) \in \mathcal{H}(\bar{F})$ but $\mathcal{H}(G_j) \oplus_{\text{h}} \mathcal{H}(H_{i,j}) \notin \mathcal{H}(\bar{F})$. Hence, since $\mathcal{H}(G_i) \oplus_{\text{h}} \mathcal{H}(H_{i,j}) = \mathcal{H}(G_i \oplus_{\text{c}} H_{i,j})$, we have $G_i \oplus_{\text{c}} H_{i,j} \in \bar{F}$. However, with the same reasoning, $G_j \oplus_{\text{c}} H_{i,j} \notin \bar{F}$. Hence, $\sim_{\bar{F}}$ has infinite index over $\bar{F}$. ☐

We can use Theorem 3 to constructively derive algorithms for deciding properties of hypergraphs of incidence treewidth at most t: assume that we have a family of t-boundaried hypergraphs $\mathcal{F} \subseteq \mathcal{H}_t^{\text{large}}$ with incidence treewidth at most t and such that $\sim_{\mathcal{F}}$ has finite index over $\mathcal{H}_t^{\text{large}}$. There is a set $F \subseteq \mathcal{U}_{t,2}^{\text{large}}$ for which $\mathcal{F} = \mathcal{H}(F)$.

We can decide in linear time whether a given hypergraph $H \in \mathcal{H}_t^{\text{large}}$ is isomorphic to a hypergraph in $\mathcal{F}$: compute the incidence graph G of H, that is, $\mathcal{H}(G) = H$. Since the graph G has treewidth at most t, we can compute a tree decomposition for G in linear time [3]. In the same way as in [7, Theorem 6.72], this tree decomposition can be converted in linear time into a parse tree T over the operators in Definition 2 that generates a graph isomorphic to G. By Theorem 3, a finite tree automaton can check whether T generates a graph G' isomorphic to a graph in F, which is the case if and only if $\mathcal{H}(G')$ is isomorphic to some graph in $\mathcal{F}$. The finite tree automaton can be constructed from the equivalence classes of $\sim_{\mathcal{F}}$.

5 Hypergraph Cutwidth is FPT

In this section we show that HYPERGRAPH CUTWIDTH is fixed-parameter linear. By Lemma 1, this implies that TRELLIS WIDTH is fixed-parameter linear as well.

Let k-HCW denote the class of all hypergraphs with cutwidth at most k. We will use Theorem 3 to show that the parse trees for graphs in k-HCW can be recognized by a finite tree automaton. To make Theorem 3 applicable, we first show that, for the hypergraphs in k-HCW, we can find a constant upper bound t on their incidence treewidth. This implies that isomorphic graphs can be generated by linear-size parse trees over the operator set in Definition 2 or, in terms of Theorem 3, that the graphs in k-HCW are isomorphic to the graphs in a subset of $\subseteq \mathcal{H}(\mathcal{U}_{t,2}^{\text{small}})$.

Lemma 3. Let $H = (V, E)$ be a hypergraph. If H has cutwidth at most k, then H has incidence treewidth at most $\max\{k, 1\}$.

Proof. Suppose H has cutwidth at most k. Let $H' = (V, E')$ denote the hypergraph obtained from H by removing all hyperedges of size at most 1. Consider a linear layout l of cutwidth at most k of the vertices of H'. Without loss of generality, assume that l maps to the natural numbers 1 to n and let $V = \{v_1, \ldots, v_n\}$ be such that $l(v_i) = i$. We construct a path decomposition for the incidence graph $\mathcal{I}(H')$ with the bags $L_1, R_1, L_2, R_2, \ldots, L_{n-1}, R_{n-1}$ and a path connecting the bags in this order. For every $i, 1 \leq i \leq n-1$, let $L_i := \theta_{l,H'}(i) \cup \{v_i\}$ and $R_i := \theta_{l,H'}(i) \cup \{v_{i+1}\}$, recalling that $\theta_{l,H'}(i)$ is the set of hyperedges that cross position i.

We now prove that this is a path decomposition for $\mathcal{I}(H')$. Let $\{v_i, e\}$ be any edge in the incidence graph of H' with $v_i \in V$ and $e \in E'$. Since $|e| \geq 2$, we have that either $e \in \theta_{l,H'}(i-1)$ or $e \in \theta_{l,H'}(i)$. Therefore, either $e \in R_{i-1}$ or $e \in L_i$. But $v_i \in R_{i-1} \cap L_i$. Thus, v_i and e occur together

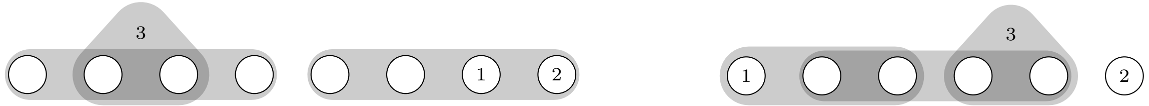

(a) A 3-boundaried hypergraph $\mathcal{H}(G)$.

(b) A 3-boundaried hypergraph $\mathcal{H}(H)$.

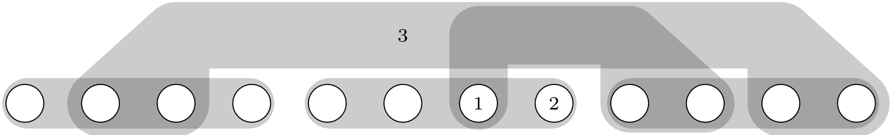

(c) The glued hypergraph $\mathcal{H}(G) \oplus_{\mathrm{h}} \mathcal{H}(H) = \mathcal{H}(G \oplus_{\mathrm{c}} H)$.

Figure 2: The two hypergraphs represented by the well-colored 3-boundaried graphs G and H in Figure 1 and the glued hypergraph $\mathcal{H}(G) \oplus_{\mathrm{h}} \mathcal{H}(H)$, which is equivalent to $\mathcal{H}(G \oplus_{\mathrm{c}} H)$.

in at least one bag. Next, consider the vertex $v_i \in V$. It occurs in the two bags R_{i-1} and L_i, which are consecutive and thus induce a connected path. Finally, consider a hyperedge $e \in E'$. It occurs in all bags $L_i, R_i, \ldots, L_j, R_j$ where v_i is the first vertex in the layout l occurring in e and v_j is the last vertex in l occurring in e. But these bags are all consecutive on the path and thus induce a connected path.

The width of this path decomposition is, by construction, $\max_{1 \leq i \leq n-1}\{|\theta_{H',l}(i)|\} \leq k$. To obtain a tree decomposition for H from this path decomposition for H', we only need to take care of hyperedges of size at most 1. For every hyperedge $e \in E$ of size 1, add a new bag $\{e, v\}$, where v is the unique vertex contained in e, and make it adjacent to an arbitrary bag containing v. For every hyperedge $e \in E$ of size 0, add a new bag $\{e\}$, and make it adjacent to an arbitrary bag. In this way, we obtain a tree decomposition for the incidence graph of H of width at most $\max\{k, 1\}$. Thus, H has incidence treewidth at most $\max\{k, 1\}$. ☐

Having Lemma 3 and Theorem 3 at our hands, it now remains to prove that the canonical right congruence $\sim_{k\text{-HCW}}$ of k-HCW has finite index. This finally shows that k-HCW can be recognized in linear time and, therefore, that HYPERGRAPH CUTWIDTH is fixed-parameter linear.

To show that $\sim_{k\text{-HCW}}$ has finite index, we show that, given a t-boundaried hypergraph G, only a finite number of bits of information about a t-boundaried hypergraph H is needed in order to decide whether $G \oplus_{\mathrm{h}} H \in k\text{-HCW}$. To this end, we employ the Method of Test Sets [7]: let $\mathcal{T}$ be a set of objects[3] called *tests*. A t-boundaried graph can *pass* a test. For t-boundaried hypergraphs G_1 and G_2, let $G_1 \sim_{\mathcal{T}} G_2$ if and only if G_1 and G_2 pass the same subset of tests in $\mathcal{T}$. Obviously, $\sim_{\mathcal{T}}$ is an equivalence relation. Our aim is to find a set $\mathcal{T}$ of tests such that $\sim_{\mathcal{T}}$ refines $\sim_{k\text{-HCW}}$ (that is, if $G_1 \sim_{\mathcal{T}} G_2$, then $G_1 \sim_{k\text{-HCW}} G_2$). This will imply that, if $\sim_{\mathcal{T}}$ has finite index, so has $\sim_{k\text{-HCW}}$. To show that $\sim_{\mathcal{T}}$ has finite index, we show that we can find a *finite* set $\mathcal{T}$ such that $\sim_{\mathcal{T}}$ refines $\sim_{k\text{-HCW}}$.

Intuitively, in our case, a hypergraph G will pass a test T if it has a restricted linear layout, where each of its boundary vertices gets mapped to predefined integer values and each of the remaining vertices “lands” within one of a set of given “landing zones” between the integer values of the real line. This restricted linear layout will impose the same restrictions on an optimal cutwidth layout for G that are also imposed by an optimal cutwidth layout of $G \oplus_{\mathrm{h}} H$ for some hypergraph H that corresponds to T.

Definition 7. Let G and H be t-boundaried hypergraphs such that $G \oplus_{\mathrm{h}} H \in k\text{-HCW}$. Let v_i denote the vertex that is mapped to position i in an optimal cutwidth layout that maps to integer values.

A *landing zone* is a tuple in $\{0, \ldots, k\} \times 2^{\{1,\ldots,t\}}$. A *size-$n$ test* $T = (\pi, S)$ consists of a map $\pi \colon \{1, \ldots, t\} \to \{1, \ldots, n\}$ and a sequence $S = (S_0, S_1, \ldots, S_n)$ of landing zones.

We define an *H-test* $T = (\pi, S)$ as follows: for each vertex $v_i \in \partial(H)$, set $\pi(\ell) := i$, where ℓ is the label of v_i. For $i \in \{0, \ldots, n\}$, $S_i := (w_i, E_i)$, where

i) w_i is the number of hyperedges in H and that contain vertices in $\{v_1, \ldots, v_i\} \cap V(H)$ and $\{v_{i+1}, \ldots, v_n\} \cap V(H)$.

ii) E_i is the set of labels of hyperedges in H containing vertices in $\{v_1, \ldots, v_i\} \cap V(H)$ and $\{v_{i+1}, \ldots, v_n\} \cap V(H)$.

Figure 3 illustrates this definition. We now formally define what it means to pass a test. The definition models the aim that if a graph G passes an H-test, then $G \oplus_{\mathrm{h}} H \in k\text{-HCW}$.

Definition 8. Let $G = (V, E)$ be a t-boundaried hypergraph and $T = (\pi, S)$ be an H-test for some t-boundaried hypergraph H, where $S = (S_0, \ldots, S_n)$ and $S_i = (w_i, E_i)$.

A *T-compatible layout for G* is an injective function $f \colon V \to \mathbb{R}$ such that each vertex $v \in \partial(G)$ with label ℓ is mapped to $\pi(\ell)$ and such that every vertex $v \in V \setminus \partial(G)$ is mapped into some open interval $(i, i+1)$ for $0 \leq i \leq n$.

The *weighted cutwidth of f* is $\max_{i \in \mathbb{R}}(|\theta_f(i)| + w_{\lfloor i \rfloor})$, where $\theta_f(i)$ is the set of hyperedges containing two vertices v and w with $f(v) < i < f(w)$ and that do not have a label in $E_{\lfloor i \rfloor}$.

Finally, G *passes the test* T if there is a T-compatible layout f for G whose weighted cutwidth is at most k.

Lemma 4. For $\mathcal{T}$ being the set of all tests, the equivalence relation $\sim_{\mathcal{T}}$ refines $\sim_{k\text{-HCW}}$.

Proof. We show that if two t-boundaried hypergraphs G_1, G_2 pass the same subset of tests of $\mathcal{T}$, then, for all t-boundaried hypergraphs H, $G_1 \oplus_{\mathrm{h}} H \in k\text{-HCW}$ if and

---

[3] For the moment, it is not important what exactly a test is.

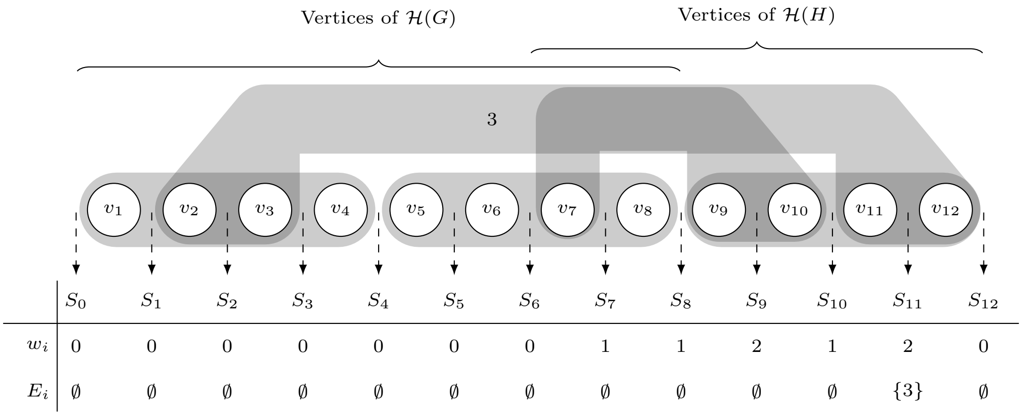

	S_0	S_1	S_2	S_3	S_4	S_5	S_6	S_7	S_8	S_9	S_{10}	S_{11}	S_{12}
w_i	0	0	0	0	0	0	0	1	1	2	1	2	0
E_i	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	$\{3\}$	$\emptyset$

Figure 3: Construction of the H-test illustrated using the glued hypergraph $\mathcal{H}(G) \oplus_h \mathcal{H}(H)$ from Figure 2.

only if $G_2 \oplus_h H \in k$-HCW. The proof is based on the following two claims that are proved afterwards.

i) If $G_1 \oplus_h H \in k$-HCW, then G_1 passes the H-test.

ii) If G_2 passes the H-test, then $G_2 \oplus_h H \in k$-HCW.

Let H be a t-boundaried hypergraph such that $G_1 \oplus_h H \in k$-HCW, and let T be an H-test. By (i), G_1 passes T. Since G_1 and G_2 pass the same tests, also G_2 passes T. By (ii), it follows that $G_2 \oplus_h H \in k$-HCW. The reverse direction is proved symmetrically. It only remains to prove (i) and (ii).

(i) Let l be a layout of $G_1 \oplus_h H$ with cutwidth at most k that only maps boundary vertices to integral positions and let T be the H-test. The restriction f of l to the vertex set of G_1 is then clearly T-compatible. Moreover, f has weighted cutwidth at most k: to this end, for an arbitrary $i \in \mathbb{N}$, we show that $|\theta_f(i)| + w_{\lfloor i \rfloor}$ from Definition 8 is at most k. Consider the set A of hyperedges of $G_1 \oplus_h H$ that contain two vertices v, w with $f(v) < i < f(w)$. Obviously, $|A| \leq k$. We show $|\theta_f(i)| + w_{\lfloor i \rfloor} = |A|$. The set A can be partitioned as $A = A^H \uplus A^{G_1} \uplus A^*$, where A^H are hyperedges in H, A^{G_1} are hyperedges in G_1 and A^* are hyperedges that are neither in H nor in G_1. The hyperedges in A^* are labeled edges that are the result of taking the unions of equally-labeled hyperedges from G_1 and H. We further partition $A^* = A^+ \uplus A^-$ such that A^+ is the set of hyperedges with labels in $E_{\lfloor i \rfloor}$. Now, by definition, we have $w_{\lfloor i \rfloor} = |A^H| + |A^+|$ and $|\theta_f(i)| = |A^{G_1}| + |A^-|$. Therefore, $|\theta_f(i)| + w_{\lfloor i \rfloor} = |A| \leq k$.

Note here, that, if we were to define $\theta_f(i)$ in Definition 8 to contain *all* edges that contain vertices left and right of i (rather than only those with no label in $E_{\lfloor i \rfloor}$), then we would obtain $|\theta_f(i)| = |A^{G_1}| + |A^+| + |A^-|$ here, that is, $|\theta_f(i)| + w_{\lfloor i \rfloor}$ would count the edges in A^+ twice.

(ii) The H-test T was obtained from a linear layout l of cutwidth k for $G_1 \oplus_h H$. Moreover, there is a T-compatible layout f for G_2. First note that l and f agree on the layout of vertices in $\partial(G_2)$ and $\partial(H)$ and that, apart from these, f lays out vertices at non-integral positions, whereas l lays out vertices at integral positions. Because of this, in a layout g for $G_2 \oplus_h H$ that lays out vertices v of H at position $l(v)$ and vertices v of G_2 at position $f(v)$, every two vertices in $G_2 \oplus_h H$ are laid out at distinct positions by g. Hence, g is injective and, therefore, a layout.

Now it is easy to observe that g is a layout of cutwidth at most k for $G_2 \oplus_h H$: to this end, for a position $i \in \mathbb{R}$, consider the set A of hyperedges of $G_2 \oplus_h H$ containing vertices v, w with $g(v) < i < g(w)$ and let it be partitioned $A = A^H \uplus A^{G_2} \uplus A^+ \uplus A^-$ in the same way as above. Again, we have $w_{\lfloor i \rfloor} = |A^H| + |A^+|$ and $|\theta_f(i)| = |A^{G_2}| + |A^-|$, yielding $|A| = |A^{G_2}| + |A^H| + |A^+| + |A^-| = \theta_f(i) + w_{\lfloor i \rfloor}$. Since G_2 passes the H-test, this is at most k. ☐

We have shown a set of tests $\mathcal{T}$ such that $\sim_{\mathcal{T}}$ refines $\sim_{k\text{-HCW}}$. However, the set $\mathcal{T}$ is infinite and, therefore, does not yet yield that $\sim_{k\text{-HCW}}$ has finite index. However, we can achieve this using the following lemma.

Lemma 5. Let G and H be t-boundaried hypergraphs. For every H-test T_1, there is a test T_2 of size $2t(2k+2)$ such that G passes T_1 if and only if it passes T_2.

Proof. Let the sequence S of landing zones of T_1 be $S_0 = (E_0, w_0), \ldots, (E_n, w_n) = S_n$. For $E \subseteq \{1, \ldots, t\}$, we call a maximal subsequence $(E_i, w_i), \ldots, (E_j, w_j)$ of S with $E_i = \cdots = E_j$ a *strait*. We first show that there are at most $2t$ straits, and then show that we can shorten each strait to a length of $2k+2$ by removing some landing zones without changing the satisfiability of the test.

For $\ell \subseteq \{1, \ldots, t\}$, let $I_\ell := \{i \leq n : \ell \in E_i\}$. Observe that I_ℓ is a consecutive subset of $\{0, \ldots, n\}$: this is by the construction of the E_i in Definition 8 from an optimal layout for some t-boundaried graph $G' \oplus_h H$. That is, each I_ℓ for some $\ell \in \{1, \ldots, t\}$ is an interval of the natural numbers with a minimum element and a maximum element, which we both call *events*. Hence, the I_ℓ for all $\ell \in \{1, \ldots, t\}$ in total have at most $2t$ events. Since straits can only start at an event, and since only one strait can start at a fixed event, it follows that S is partitioned into at most $2t$ straits.

It remains to shorten the straits. Let $(E_i, w_i), \ldots, (E_j, w_j)$ be a strait. We apply the data reduction rules (R1–R3) from the proof of Theorem 6.82 by Downey and Fellows [7], which are based on the following observations: deleting from S one of two consecutive landing zones $(E_i, w_i), (E_{i+1}, w_{i+1})$ with $E_i = E_{i+1}$ and $w_i = w_{i+1}$ yields an equivalent test, as does adding such landing zones. Moreover, if we replace a landing zone (E_i, w_i) by a landing zone (E_i, w_i') with $w_i' \leq w_i$, we obtain a test that is easier to pass.

Since the w_i are, by construction of T_1, bounded from above by k, we can transform T_1 into T_2 by these data reduction rules so that each of the $2t$ straits has length at most $(2k+2)$ [7, Theorem 6.83]. $\square$

Now the following theorem is easy to prove.

Theorem 4. Hypergraph Cutwidth *is fixed-parameter linear.*

Proof. Lemma 3 shows that graphs in k-HCW have constant treewidth at most t, and therefore, that all such hypergraphs can be linear-time transformed into linear-size parse trees for t-boundaried hypergraphs. Using Theorem 3, we show that parse trees corresponding to hypergraphs in k-HCW are recognizable by a finite tree automaton: let $\mathcal{T}$ be the set of all tests and $\mathcal{T}'$ be the set of all tests of size $2t(2k+2)$. Lemma 5 shows that $\sim_{\mathcal{T}'}$ refines $\sim_{\mathcal{T}}$. Lemma 4 shows that $\sim_{\mathcal{T}}$ refines $\sim_{k\text{-HCW}}$. Therefore, $\sim_{k\text{-HCW}}$ has at most as many equivalence classes as $\sim_{\mathcal{T}'}$. Since k and t are constant, $\mathcal{T}'$ is finite, implying finite index for $\sim_{\mathcal{T}'}$ and, consequently, for $\sim_{k\text{-HCW}}$. $\square$

6 Hypertree Width and Variants

In this section we show a negative application of our hypergraph Myhill-Nerode analog to Hypertree Width, Generalized Hypertree Width, and Fractional Hypertree Width. The problems are, given a hypergraph H and an integer k as input, to decide whether H has (generalized, fractional) hypertree width at most k. Since Hypertree Width is W[2]-hard [19], Generalized Hypertree Width is paraNP-hard [20] (NP-hard for $k=3$), and the parameterized complexity of Fractional Hypertree Width is open but expected to be paraNP-hard [26], it would be nice to find non-standard parameters for which these problems are fixed-parameter tractable. While it is known that the generalized hypertree width of a hypergraph is at most the incidence treewidth plus one [11], the incidence treewidth may be arbitrarily large even for hypergraphs with (generalized, fractional) hypertree width one, since adding a universal hyperedge to any hypergraph reduces its (generalized, fractional) hypertree width to one. This makes incidence treewidth a natural parameter to explore. However, the following theorem shows that Hypertree Width, Generalized Hypertree Width, and Fractional Hypertree Width do not have finite index. Moreover, the construction we use in the proof leads us to believe that the problems are W[1]-hard with respect to the parameter incidence treewidth.

Theorem 5. *Let $k \geq 0$. Let HTW_k (respectively, GHTW_k and FHTW_k) be the family of all incidence graphs G such that $\mathcal{H}(G)$ has hypertree width (respectively, generalized hypertree width and fractional hypertree width) at most k. The canonical right congruences $\sim_{\mathcal{H}(\mathsf{HTW}_k)}$, $\sim_{\mathcal{H}(\mathsf{GHTW}_k)}$, and $\sim_{\mathcal{H}(\mathsf{FHTW}_k)}$ do not have finite index over $\mathcal{H}(\mathcal{U}_{t,2}^{small})$.*

Proof Outline. In this proof outline, we give a construction of a t-boundaried hypergraph H_n with bounded incidence treewidth, for every $n \geq 1$. Then we outline how one can show that $H_n \oplus_{\mathrm{h}} H_m$ has (generalized, fractional) hypertree width 4 if and only if $n = m$. This implies an infinite number of equivalence classes for the canonical right congruences $\sim_{\mathcal{H}(\mathsf{HTW}_4)}$ and $\sim_{\mathcal{H}(\mathsf{GHTW}_4)}$. We refer to Appendix A for a detailed proof.

For every $n \geq 1$, we construct a t-boundaried hypergraph H_n with $t = 28$, generalized hypertree width 4, and incidence treewidth at most 12. The vertex set of H_n is $V = A \cup B \cup C \cup D \cup S \cup T \cup X$, where $A = \{a, y\}, B = \{b, z\}, C = \{c, y\}, D = \{d, z\}, S = \{s_1, \ldots, s_8\}, T = \{t_1, \ldots, t_8\}$ and $X = \{x_1, \ldots, x_{6n}\}$. The hyperedge set of H_n is $E = \{A, B, C, D\} \cup B_S \cup \{S_c, S_d, S_y, S_z\} \cup B_T \cup \{T_a, T_b, T_y, T_z\} \cup \{E_{3i}, E_{3i+1} : 1 \leq i < 2n\} \cup \{E_{i,i+1} : 1 \leq i < 6n\}$, where

$$
\begin{aligned}
&B_S \text{ is the set of all possible binary hyperedges on } S,\\
&S_c = \{c, s_1, s_2\}, S_d = \{d, s_3, s_4\},\\
&S_y = \{y, s_5, s_6\}, S_z = \{z, s_7, s_8\}\\
&B_T \text{ is the set of all possible binary hyperedges on } T,\\
&T_a = \{a, t_1, t_2\}, T_b = \{b, t_3, t_4\},\\
&T_y = \{y, t_5, t_6\}, T_z = \{z, t_7, t_8\}\\
E_1 &= \{s_8, x_1\},\\
E_{3i} &= \{a, c, y, x_{3i}\} \text{ for } 1 \leq i < 2n,\\
E_{3i+1} &= \{b, d, z, x_{3i+1}\} \text{ for } 1 \leq i < 2n,\\
E_{6n} &= \{x_{6n}, t_1\},\\
E_{6i+1,6i+2} &= \{a, b, x_{6i+1}, x_{6i+2}\} \text{ for } 0 \leq i < n,\\
E_{6i+4,6i+5} &= \{c, d, x_{6i+4}, x_{6i+5}\} \text{ for } 0 \leq i < n, \text{ and}\\
E_{3i,3i+1} &= \{x_{3i}, x_{3i+1}\} \text{ for } 1 \leq i < 2n.
\end{aligned}
$$

The set of boundary hyperedges is $\{A, B, C, D, S_c, S_d, S_y, S_z, T_a, T_b, T_y, T_z\}$. The set of boundary vertices is $S \cup T$. They are labeled from 1 to 28 in this order and by increasing indices. See Figure 4 for an illustration of H_2 induced on $V \setminus (S \cup T)$.

Consider a tree decomposition for $H_n \oplus_{\mathrm{h}} H_m$ with (generalized) hypertree width 4. The vertex sets S and T and the hyperedges containing them make sure that some bag $\mathcal{B}_S$ of the decomposition contains all of S and that some bag $\mathcal{B}_T$ contains all of T. Now, both in $\mathcal{G}(H_n)$ and in $\mathcal{G}(H_m)$, there is a path from a vertex in S to a vertex in T passing through all vertices x_i by increasing indices. The edges of this path are covered by intermediate bags lying on the path from $\mathcal{B}_S$ to $\mathcal{B}_T$ on the tree decomposition. Observe that no vertex x_i is contained in a boundary hyperedge. Therefore, when we restrict the tree decomposition to the vertices in H_n, we recover a tree decomposition for H_n where all intermediate bags are covered by at most three hyperedges. Moreover, our construction makes sure that when a bag is covered by three hyperedges, at least two of them are boundary hyperedges. In every such tree decomposition for H_n, when considering the intermediate bags starting from $\mathcal{B}_S$ that contain either A, B or C, D in their cover, we first encounter bags covered by C, D, then bags covered by A, B, then bags covered by C, D, and so on, and there are exactly n alternations from C, D to A, B in this sequence. Therefore, in order to be able to merge such decompositions for H_n and H_m, we must have $n = m$. $\square$

The number of equivalence classes observed in Theorem 5 entails a lower bound on the amount of information that

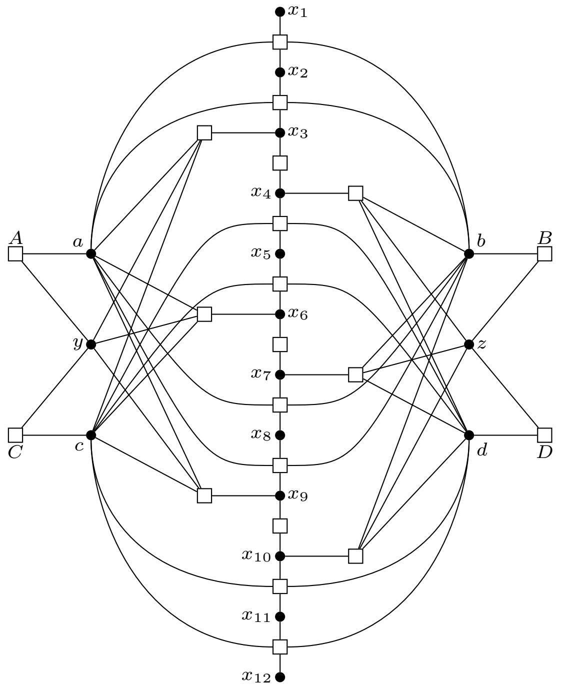

Figure 4: The incidence graph of H_2 induced on $V \setminus (S \cup T)$. Boxes represent hyperedges.

needs to be maintained by an algorithm when it decides whether a given hypergraph has (generalized, fractional) hypertree width k using a tree decomposition of the incidence graph: typical such algorithms implicitly remember for each bag of the tree decomposition (by means of a dynamic programming table or by the states of a tree automaton) which equivalence classes of the bag under consideration can complete into a graph with (generalized, fractional) hypertree width k.

However, restricting the construction in the proof of Theorem 5 to at most n vertices, the proof of Theorem 5 exhibits a class $\mathcal{C}$ of t-boundaried hypergraphs on at most n vertices with constant incidence treewidth and constant t for which the canonical right congruences have $\Omega(n)$ equivalence classes. Now, consider a class $\mathcal{C}'$ of $O(k)$-boundaried hypergraphs where each hypergraph contains k copies of hypergraphs from $\mathcal{C}$ and has at most n' vertices. Then, the number of equivalence classes of the canonical right congruences is $\Omega((n'/k)^k)$ for $\mathcal{C}'$. Hence, we conjecture that an algorithm with running time $f(k) \cdot n^c$ for a constant c and f only depending on k does not exist.

Conjecture 1. HYPERTREE WIDTH, GENERALIZED HYPERTREE WIDTH, *and* FRACTIONAL HYPERTREE WIDTH *are W[1]-hard with respect to the parameter incidence treewidth.*

7 Conclusion

We have extended the graph analog of the Myhill-Nerode theorem to colored graphs and hypergraphs, making the methodology more widely applicable. Our positive application shows that HYPERGRAPH CUTWIDTH, TRELLIS WIDTH, and MATROID PATHWIDTH are fixed-parameter linear. As a negative application, we showed that HYPERTREE WIDTH, GENERALIZED HYPERTREE WIDTH, and FRACTIONAL HYPERTREE WIDTH do not have finite index, and therefore the parse trees associated to Yes-instances of bounded incidence treewidth cannot be recognized by finite tree automata.

Acknowledgments

The authors are thankful to Dr. Mahdi Parsa from Griffith University, Brisbane, Australia, for fruitful discussions. René van Bevern acknowledges support from the DFG, project DAPA, NI 369/12. Serge Gaspers acknowledges support from the Australian Research Council (grant DE120101761), as do Michael R. Fellows and Frances A. Rosamond (grants DP 1097129 and DP 110101792).

Bibliography

[1] K. R. Abrahamson and M. R. Fellows. Finite automata, bounded treewidth, and well-quasiordering. In *Graph Structure Theory*, volume 147 of *Contemporary Mathematics*, pages 539–564. American Mathematical Society, 1991. Cited on page 3.

[2] L. Bahl, J. Cocke, F. Jelinek, and J. Raviv. Optimal decoding of linear codes for minimizing symbol error rate. *IEEE Inform. Theory*, 20(2):284 – 287, 1974. Cited on page 3.

[3] H. L. Bodlaender. A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM J. Comput.*, 25(6):1305–1317, 1996. Cited on pages 4 and 6.

[4] H. L. Bodlaender, M. R. Fellows, and D. M. Thilikos. Derivation of algorithms for cutwidth and related graph layout parameters. *J. Comput. Syst. Sci.*, 75(4):231–244, 2009. Cited on page 3.

[5] J. Cahoon and S. Sahni. Exact algorithms for special cases of the board permutation problem. In *Proceedings of the 21st Annual Allerton Conference on Communication, Control, and Computing*, pages 246–255, 1983. Cited on page 3.

[6] R. Diestel. *Graph Theory*, volume 173 of *Graduate Texts in Mathematics*. Springer Verlag, New York, 4th edition, 2010. Cited on page 2.

[7] R. G. Downey and M. R. Fellows. *Parameterized Complexity*. Springer, 1999. Cited on pages 2, 3, 4, 5, 6, 7, 8, and 9.

[8] M. R. Fellows and M. A. Langston. An analogue of the Myhill-Nerode Theorem and its use in computing finite-basis characterizations (extended abstract). In *Proc. 30th FOCS*, pages 520–525. IEEE Computer Society, 1989. Cited on pages 2 and 4.

[9] M. R. Fellows and M. A. Langston. On well-partial-order theory and its application to combinatorial problems of VLSI design. *SIAM J. Discrete Math.*, 5(1):117–126, 1992. Cited on page 3.

[10] M. R. Fellows and M. A. Langston. On search, decision, and the efficiency of polynomial-time algorithms. *J. Comput. Syst. Sci.*, 49(3):769–779, 1994. Cited on pages 3 and 4.

[11] F. V. Fomin, P. A. Golovach, and D. M. Thilikos. Approximating acyclicity parameters of sparse hypergraphs. In *Proc. 26th STACS*, volume 3 of *LIPIcs*, pages 445–456. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Germany, 2009. Cited on page 9.

[12] G. D. Forney, Jr. Final report on a coding system designed for advanced solar missions. Technical Report NAS2-3637, NASA Ames Research Center, 1967. Cited on page 1.

[13] G. D. Forney, Jr. Coset codes. II. binary lattices and related codes. *IEEE Inform. Theory*, 34(5):1152 –1187, 1988. Cited on page 3.

[14] G. D. Forney, Jr. Dimension/length profiles and trellis complexity of linear block codes. *IEEE Inform. Theory*, 40(6):1741–1752, 1994. Cited on page 4.

[15] G. D. Forney, Jr. Codes on graphs: constraint complexity of cycle-free realizations of linear codes. *IEEE Inform. Theory*, 49(7):1597–1610, 2003. Cited on page 1.

[16] G. D. Forney, Jr. and M. D. Trott. The dynamics of group codes: State spaces, trellis diagrams and canonical encoders. *IEEE Inform. Theory*, 39:1491–1513, 1993. Cited on page 3.

[17] F. Gavril. Some NP-complete problems on graphs. In *Proc. 1977 Conf. on Inf. Sc. and Syst.*, pages 91–95. Johns Hopkins Univ., 1977. Cited on page 3.

[18] J. F. Geelen, A. M. H. Gerards, and G. P. Whittle. Towards a matroid-minor structure theory. In *Combinatorics, Complexity, and Chance: A Tribute to Dominic Welsh*, number 34 in Oxford Lecture Series in Mathematics and its Applications, pages 72–82. Oxford University Press, 2007. Cited on page 1.

[19] G. Gottlob, M. Grohe, N. Musliu, M. Samer, and F. Scarcello. Hypertree decompositions: Structure, algorithms, and applications. In *Proc. 31st WG*, volume 3787 of *Lecture Notes in Computer Science*, pages 1–15. Springer, 2005. Cited on page 9.

[20] G. Gottlob, Z. Miklós, and T. Schwentick. Generalized hypertree decompositions: NP-hardness and tractable variants. *J. ACM*, 56(6), 2009. Cited on page 9.

[21] P. Hliněný. On matroid properties definable in the MSO logic. In *Proc. 28th MFCS*, volume 2747 of *Lecture Notes in Computer Science*, pages 470–479. Springer, 2003. Cited on page 1.

[22] N. Kashyap. The "art of trellis decoding" is NP-hard. In *Proc. 17th AAECC*, AAECC'07, pages 198–207. Springer, 2007. Cited on pages 1 and 3.

[23] N. Kashyap. Matroid pathwidth and code trellis complexity. *SIAM J. Discret. Math.*, 22(1):256–272, 2008. Cited on pages 1, 2, and 3.

[24] P. G. Kolaitis and M. Y. Vardi. Conjunctive-query containment and constraint satisfaction. *J. Comput. Sci.*, 61 (2):302–332, 2000. Cited on page 2.

[25] F. R. Kschischang and V. Sorokine. On the trellis structure of block codes. *IEEE Inform. Theory*, 41(6):1924–1937, 1995. Cited on page 3.

[26] D. Marx. Approximating fractional hypertree width. *ACM Transactions on Algorithms*, 6(2), 2010. Cited on page 9.

[27] J. Massey. Foundation and methods of channel encoding. In *Proc. Int. Conf. Information Theory and Systems*, volume 65 of *NTG-Fachberichte*, Berlin, Germany, 1978. Cited on pages 1 and 3.

[28] Z. Miller and I. H. Sudborough. A polynomial algorithm for recognizing bounded cutwidth in hypergraphs. *Mathematical Systems Theory*, 24(1):11–40, 1991. Cited on pages 2 and 3.

[29] D. Muder. Minimal trellises for block codes. *IEEE Inform. Theory*, 34(5):1049–1053, 1988. Cited on page 3.

[30] M. Samer and S. Szeider. Constraint satisfaction with bounded treewidth revisited. *J. Comput. Sci.*, 76(2):103–114, 2010. Cited on page 2.

[31] D. M. Thilikos, M. J. Serna, and H. L. Bodlaender. Cutwidth I: A linear time fixed parameter algorithm. *J. Algorithms*, 56(1):1–24, 2005. Cited on page 3.

[32] A. Vardy. Trellis structure of codes. In *Handbook of Coding Theory*, chapter 24. Elsevier, 1998. Cited on pages 1 and 3.

[33] A. J. Viterbi. Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Inform. Theory*, 13(2):260–269, 1967. Cited on pages 1 and 3.

A Proof of Theorem 5

We continue the proof of Theorem 5. Recall the construction of H_n from the proof outline.

In this construction, the vertices in S and T and the hyperedges containing them are only used to make sure that every tree decomposition of H_n with (generalized, fractional) hypertree width 4 contains a bag $\mathcal{B}_{-1}$ with the vertices $S \cup \{c, d, y\}$ and a bag $\mathcal{B}_{6n+1}$ with the vertices $T \cup \{b, y, z\}$. Since the sets $S \cup \{c, d, y, z\}$ and $T \cup \{a, b, y, z\}$ can also be covered by four hyperedges, all of which are boundary hyperedges, let $\mathcal{D} = (\{X_i : i \in I\}, T)$ be a tree decomposition for H_n with the bags $\mathcal{B}_{-1} = S \cup \{c, d, y, z\}$ and $\mathcal{B}_{6n+1} = T \cup \{a, b, y, z\}$. We observe that all other vertices of H_n occur in bags that are in the same connected component of the forest obtained from $\mathcal{D}$ by removing these two bags.

Claim 1. The tree decomposition $\mathcal{D}$ contains a bag $\mathcal{B}_i$, $0 \leq i \leq 6n$, with $\{s_8, x_1, c, d, y, z\} \subseteq B_0$, $\{t_1, x_{6n}, a, b, y, z\} \subseteq B_{6n}$, and $\{a, b, c, d, y, z, x_i, x_{i+1}\} \subseteq \mathcal{B}_i$, for every i, $1 \leq i < 6n$.

Proof. The primal graph $\mathcal{G}(H_n)$ contains the cliques $\{s_8, x_1\}$, $\{a, b, x_1, x_2\}$, $\{a, b, x_2, x_3\}$, $\{a, c, y, x_3\}$, $\{x_3, x_4\}$, $\{x_4, b, d, z\}$, $\{c, d, x_4, x_5\}$, $\{c, d, x_5, x_6\}$, $\{a, c, y, x_6\}$, $\{x_6, x_7\}$, $\{b, d, z, x_7\}$, $\{a, b, x_7, x_8\}$, $\ldots$, $\{x_{6n}, t_1\}$, and every two consecutive cliques in this list intersect in at least one vertex. In particular, we observe the path $(s_8, x_1, x_2, \ldots, x_{6n}, t_1)$ in $\mathcal{G}(H_n)$. Thus, $\mathcal{D}$ contains bags $\mathcal{B}_0 \supseteq \{s_8, x_1\}$, $\mathcal{B}_{6n} \supseteq \{t_1, x_{6n}\}$, and $\mathcal{B}_i \supseteq \{x_i, x_{i+1}\}$, $1 \leq i < 6n$. Moreover, each $\mathcal{B}_i$, $0 \leq i < 6n$, contains c, d, y, z since $\mathcal{B}_{-1}$ contains c, d, y, z, $\mathcal{B}_{6n-1}$ contains c, d, $\mathcal{B}_{6n+1}$ contains y, z, and without loss of generality, we can assume the $\mathcal{B}_i$, $0 \leq i \leq 6n$, were chosen such that they are on the path from $\mathcal{B}_{-1}$ to $\mathcal{B}_{6n+1}$ in T. Similarly, each $\mathcal{B}_i$, $1 \leq i \leq 6n$, contains a, b. ☐

A tree decomposition for H_n is a *good* tree decomposition if it contains the bags $\mathcal{B}_{-1} = S \cup \{c, d, y, z\}$ and $\mathcal{B}_{6n+1} = T \cup \{a, b, y, z\}$ and every bag except $\mathcal{B}_{-1}$ and $\mathcal{B}_{6n+1}$ can be covered with at most three hyperedges, and in case such a bag is covered with exactly three hyperedges, two of these hyperedges are in the boundary. A *good* cover for a good tree decomposition is a cover for each bag according to the specifications of a good tree decomposition.

Claim 2. If $\mathcal{D}$ is a good tree decomposition for H_n, then, for every i, $-1 \leq i \leq 6n$, there is a path from the bag $\mathcal{B}_i$ to the bag $\mathcal{B}_{i+1}$ that avoids all the bags $\mathcal{B}_j$, $j \in \{-1, \ldots, 6n+1\} \setminus \{i, i+1\}$.

Proof. Suppose the path from $\mathcal{B}_i$ to $\mathcal{B}_{i+1}$ passes through $\mathcal{B}_j$ with $j \in \{-1, \ldots, 6n+1\} \setminus \{i, i+1\}$. Since every bag on the path from $\mathcal{B}_i$ to $\mathcal{B}_{i+1}$ contains $\mathcal{B}_i \cap \mathcal{B}_{i+1}$, we have that $x_{i+1} \in \mathcal{B}_j$. But then $\{s_8, x_1, c, d, y, z, x_{i+1}\} \subseteq B_j$ (if $j = 0$) or $\{t_1, x_{6n}, a, b, y, z, x_{i+1}\} \subseteq B_j$ (if $j = 6n$) or $\{a, b, c, d, y, z, x_j, x_{j+1}, x_{i+1}\} \subseteq \mathcal{B}_j$ (otherwise), implying that B_j cannot be covered by two hyperedges and it cannot be covered by three hyperedges of which two are in the boundary. ☐

Claim 3. In every good cover, $\mathcal{B}_0$ is covered by $\{E_1, C, D\}$, $\mathcal{B}_{6n}$ is covered by $\{E_{6n}, A, B\}$, and for every i, $1 \leq i < 6n$,

$$\mathcal{B}_i \text{ is covered by } \begin{cases} \{E_{i,i+1}, C, D\} & \text{if } i \equiv 1 \pmod 6, \\ \{E_{i,i+1}, C, D\} & \text{if } i \equiv 2 \pmod 6, \\ \{E_i, E_{i+1}\} & \text{if } i \equiv 3 \pmod 6, \\ \{E_{i,i+1}, A, B\} & \text{if } i \equiv 4 \pmod 6, \\ \{E_{i,i+1}, A, B\} & \text{if } i \equiv 5 \pmod 6, \\ \{E_i, E_{i+1}\} & \text{if } i \equiv 0 \pmod 6. \end{cases}$$

Proof. The claim easily follows from Claim 1. ☐

Suppose $\mathcal{D}$ is a good tree decomposition for H_n. The *backbone* of $\mathcal{D}$ is the path P in T starting at the bag $\mathcal{B}_{-1}$ and ending at the bag $\mathcal{B}_{6n+1}$. By Claim 2, P visits $\mathcal{B}_0, \mathcal{B}_1, \ldots, \mathcal{B}_{6n}$ in this order. Let $P_{i,j}$ denote the subpath of P starting at $\mathcal{B}_i$ and ending at $\mathcal{B}_j$.

Claim 4. For every $i \in \{0, 6, 12, \cdots, 6n-6\}$, no bag on $P_{i,i+3}$ is covered by a set of hyperedges $\mathcal{Q}$ with $A, B \in \mathcal{Q}$ in a good cover.

Proof. Consider a bag $\mathcal{B}$ on $P_{i,i+3}$ and let $\mathcal{Q} \supseteq \{A, B\}$ be a cover for $\mathcal{B}$. The bag $\mathcal{B}$ contains the intersection of two bags that are consecutive in the list $\mathcal{B}_i, \mathcal{B}_{i+1}, \mathcal{B}_{i+1}, \mathcal{B}_{i+3}$. Therefore, at least one of $x_{i+1}, x_{i+2}, x_{i+3}$ is in $\mathcal{B}$. We also have that $c, d \in \mathcal{B}$ since $c, d \in \mathcal{B}_i \cap \mathcal{B}_{i+3}$. However, no hyperedge contains x_{i+1}, c, d or x_{i+2}, c, d or x_{i+3}, c, d. Thus, $|\mathcal{Q}| \geq 4$, and therefore $\mathcal{Q}$ is not part of a good cover. ☐

Claim 5. For every $i \in \{3, 9, 15, \cdots, 6n-3\}$, no bag on $P_{i,i+3}$ is covered by a set of hyperedges $\mathcal{Q}$ with $C, D \in \mathcal{Q}$ in a good cover.

Proof. The proof is symmetric to the proof of Claim 4. ☐

Consider a good cover of $\mathcal{D}$. A *switch* is an inclusion-wise minimal subpath $(Y_i, \ldots, Y_j)$ of the backbone of $\mathcal{D}$ where Y_i is covered by $\mathcal{Q}_i$ with $C, D \in \mathcal{Q}_i$ and Y_j is covered by $\mathcal{Q}_j$ with $A, B \in \mathcal{Q}_j$. The *signature* of a good cover of $\mathcal{D}$ is its number of switches.

Claim 6. Each good cover of each good tree decomposition of H_n has signature n.

Proof. The claim follows from Claims 3, 4, and 5. ☐

Due to Claim 6, we can speak of the signature of H_n and the signature of a good tree decomposition of H_n as the signature of some good cover of such a tree decomposition.

Let $H = H_n$ and $H' = H_m$. Consider a tree decomposition $\mathcal{D} = (\{X_i : i \in I\}, T)$ of $H \oplus_{\mathrm{h}} H'$ with generalized hypertree width 4. Without loss of generality, suppose the bags $\mathcal{B}_{-1} = S \cup \{c, d, y, z\}$ and $\mathcal{B}_{6n+1} = T \cup \{a, b, y, z\}$ are leafs of this decomposition and their neighboring bags contain both copies of x_1 and x_{6n}, respectively. Let $\mathcal{D}_{|_H}$ denote the restriction of $\mathcal{D}$ to H, i.e., it has the same tree, but each bag is restricted to the vertices of H.

Claim 7. $\mathcal{D}_{|_H}$ is a good tree decomposition for H.

Proof. Consider a bag $\mathcal{B}$ of $\mathcal{D}$ besides $\mathcal{B}_{-1}$ and $\mathcal{B}_{6n+1}$. The bag $\mathcal{B}$ contains a copy of some x_i from H'. This vertex is covered by some hyperedge from H' that does not belong to the boundary. Therefore, $\mathcal{B}_{|_H}$ is covered by at most three hyperedges. Suppose $\mathcal{B}_{|_H}$ is covered by exactly three hyperedges. Then, the cover of $\mathcal{B}_{|_{H'}}$ contains at most one hyperedge that does not belong to the boundary. But, since each such hyperedge covers at most two vertices among $\{a, b, c, d\}$, the cover of $\mathcal{B}$ contains at least two boundary hyperedges. This proves the claim. ☐

Symmetrically, $\mathcal{D}_{|_{H'}}$ is a good tree decomposition for H'. Since $\mathcal{D}_{|_H}$ and $\mathcal{D}_{|_{H'}}$ have the same signature, we conclude that $n = m$ due to Claim 6. This proves that the canonical right congruence $\sim_{\mathcal{H}(\mathsf{GHTW}_4)}$ does not have finite index over $\mathcal{H}(\mathcal{U}_{t,2}^{\text{small}})$.

To see that the same proof applies to $\sim_{\mathcal{H}(\mathsf{HTW}_4)}$, observe that in our construction every hyperedge covering a bag is a subset of that bag. To see that it extends to $\sim_{\mathcal{H}(\mathsf{FHTW}_4)}$, observe that for every bag $\mathcal{B}_i$, $0 \leq i \leq 6n$, an optimal fractional cover is integral, and Claim 4 can be extended to $A, B \in \mathcal{Q}$ with weight 1—similarly for Claim 5. This completes the proof of Theorem 5.